

Australian Advanced Computing Facility for Climate and Earth Systems Research

Ben Evans

Joseph Antony

National Computational Infrastructure
(NCI) program
nf.nci.org.au

First Stage:
Building a Peak Compute Facility
to support Climate and Earth
System Research

**NCI – National Computational Infrastructure.
Supported by NCRIS initiative from
Australian Government (DIISR).**

- **ANU with partners CSIRO, Bureau of Met, Geoscience Australia, Research Intensive Universities (Go8).**
- **Tier-1 computing capability – Peak (ANU) plus national access to small number of specialised systems**
- **New initiative in data-intensive application support.**

Building the case for Petascale computing for *Climate and Earth Systems Research*

- Build a comprehensive facility: HPC, data management and analysis, viz, web, ...
- Interoperate between research & operations.
- Need to develop cohort of specialists that support the collaborative.
- **Decision:** Proceed with Joint procurement between Bureau of Met as operational system & National Facility as research system; increase data sharing and access— nationally and internationally.

Other Tier1 groups also to use the system.



nf.nci.org.au

National Computational Infrastructure

NCI National Facility

NCI National Facility Supercomputer Stage 1 Sept 09, Stage 2 Dec 09 Bureau NWP operation system: Nov-Dec 09




(Image of Jülich, Germany)

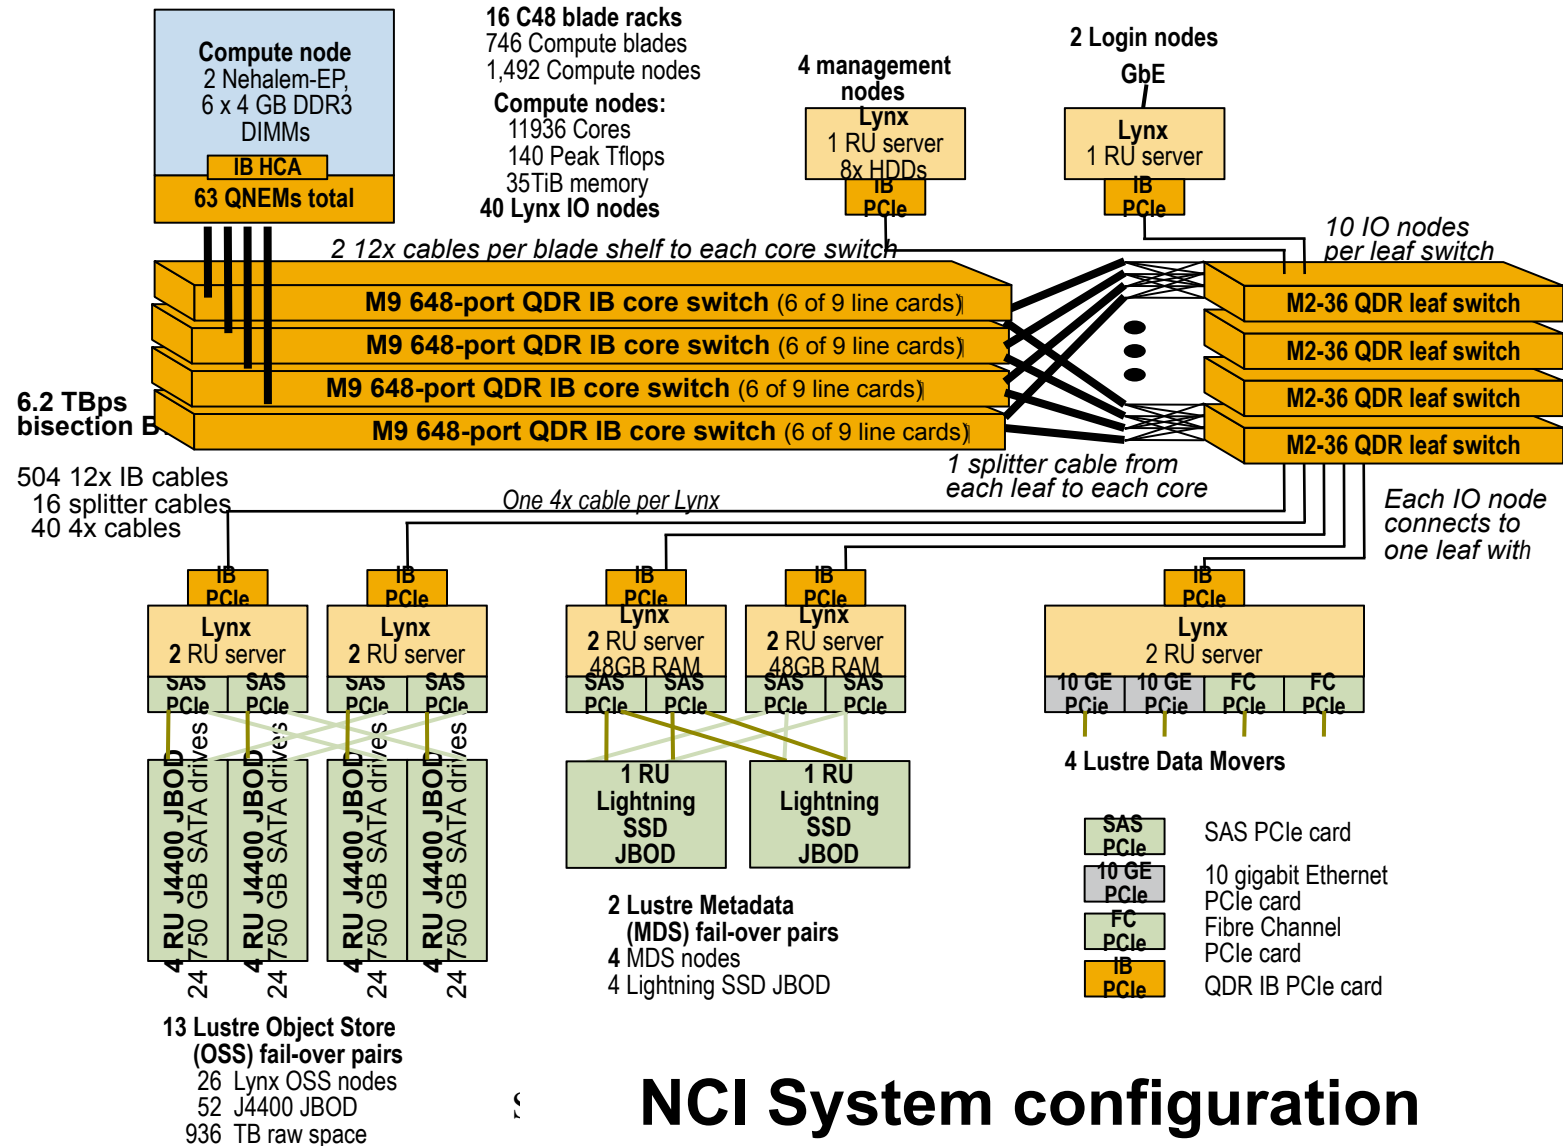
NF System details: Nodes

- 1496 nodes each with:
 - two 2.93GHz Intel Nehalem processors (quad-core)
 - 24GB of fast (DDR3-1333) memory
 - QDR (40Gb/s) Infiniband interconnect
 - 24GB flash DIMM for swap/OS
- Almost 12,000 compute cores
- 35 Tbytes of mem.

Bureau Op system

- 576 nodes (4608 cores)
- 5x performance of SX-6





NCI System configuration

Interconnect & Storage

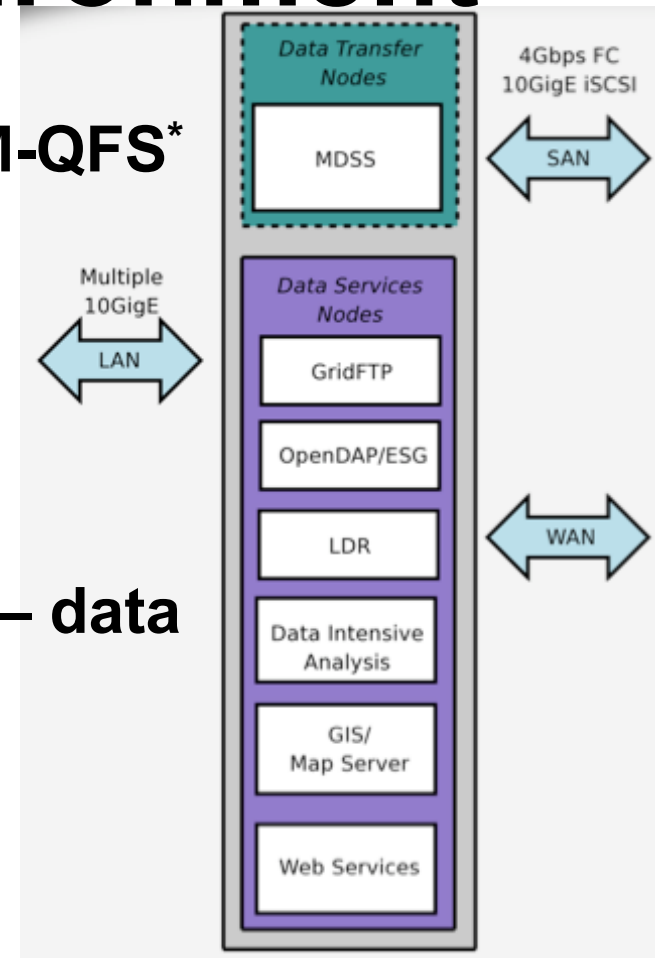
- Full QDR (40Gb/s) fat-tree Infiniband network
- Four 648-port “M9” Infiniband switches
- 2.5GB/s bandwidth and $< 2\mu\text{s}$ latency
- Lustre: Approx 1PB in raw capacity, 1248 Enterprise SATA disks
- Approx 20GB/s bandwidth




Data Intensive Environment

Capacity

- ~0.6PB of existing data under Sun SAM-QFS*
- ~3PB capacity (HSM)
- ~0.25Pbytes of online filesystems
- 30Tbytes of relational databases
- Centos (under ESX)
- Next increase disk ~2 PByte disk
- Specialised data services and hosting – data storage and transfer.
- Networks
 - 4 x 10GigE between Vayu and the DC
 - 10GigE AARNet link to SX-Transport.



Building a Facility for Climate and Earth System Research: the next Stage



National Framework for Climate Change Science (NFCCS) by the Department of Climate Change (DCC).

- to understand and predict changes in greenhouse gas levels, so that global targets to reduce emissions achieve what they are designed to do**
- to provide better information about the likely future climate at a regional level to facilitate management of Australia's increasing demand for water**
- to provide quality information about likely changes in sea level, ocean temperature, storm surge and extreme weather and climate events to facilitate management of the marine, coastal and inland environments**
- to provide the information needed for national adaptation initiatives to minimise disruption/costs associated with changes in extreme weather and climate**
- to improve our ability to predict atmospheric behaviour across short, medium long time scale**

To accomplish this:

- **Make a comprehensive service and long-range plan for climate research and earth systems.**
- **\$50M (2010 budget) to NCI to provide for “new supercomputing infrastructure to analyse and model information on climate change, earth systems and national water management”.**

Provide a comprehensive service:

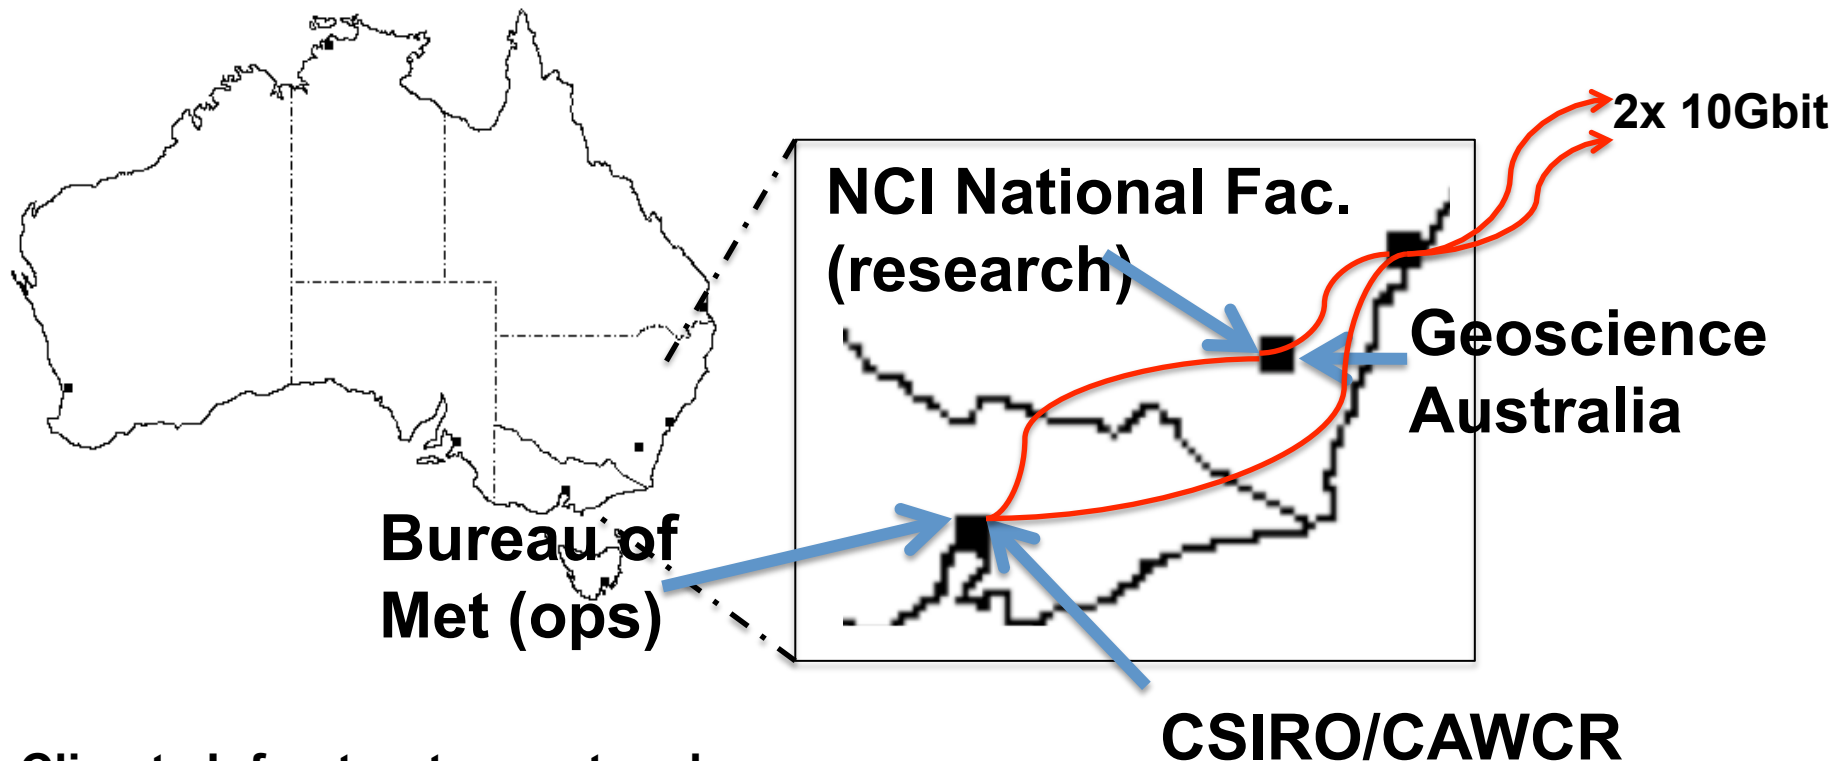
- **an increased “petascale” high performance computing facility.**
- **large scale data infrastructure and management services that provides a national data facility for research undertaken and a repository for datasets of international significance, serving also as international datanode for CMIP5 (AR5) data (ESG supernode and gateway)**
- **Build on data node for major climate collections and gateway for comprehensive climate service.**
- **facilities to provide data analysis, data synthesis, data mining and advanced visualisation.**
- **applications support and a software engineering / development team to deliver on research outcomes.**



nf.nci.org.au

National Computational Infrastructure

NCI National Facility



Climate Infrastructure network

Data transfer between NCI, and CSIRO/Bureau

CAWCR – AR5/CMIP5 run at NCI-NF (ACCESS,Mk3.6,...)

Data Storage at NCI (inc Data Core Note), Bureau and CSIRO.

Researchers (Modellers, Data Analysis, ...) connect to the NCI-NF
(UNSW, Melbourne, Monash, UTAS, ANU, CSIRO, Bureau...)

Datasets (indicative)

Main datasets

- **AR4/AR5, CMIP3/CMIP5 Australian model runs (ACCESS, and Mark 3.6)**
- **Full CMIP3 mirror**
- **NCEP R1 and R2**
- **ECMWF ERA40 + interim reanalysis**
- **CSIRO Mark2 and Mark3 models (3.0, 3.5, 3.6 – more fields than in CMIP3)**

Additional

- **CCAM model historical experiments (15+ years of data)**
- **RAMS experiments**
- **High quality Bureau rainfall and temperature (from Station data, but also produced optimally interpolated set for the country).**
- **possibly transfer realtime observations received and archived by the Bureau, etc.**



Datasets (cont)

- **WOA05**
- **HADISST**
- **Range of small, high quality but useful datasets.**
- **key model intercomparison runs**
- **Japan Meteorological Administration (JMA)**
- **ocean reanalyses forcing data used in key model intercomparison experiments**
- **observations from key special observing campaigns**
- **satellite data such as Cloudsat, TRIMM**
- **Australian Water Availability Project (AWAP)**
- **...**

Deployment of ESG Data Node and National Gateway

**Joseph Antony
and Ahmed El Zein**

ESG Installation and Modification

- Created a set of deployment notes based on a JeOS VM to improve maintainability.
- Depend on upstream O/S or service providers for Java middleware, MySQL, PostgreSQL updates
- As a result we now have distinct components for our ESG data node
 - Tomcat, PostgreSQL components
 - ESG specific components (Python modules, THREDDS)
- Next will get 'yum install esg-datanode' working

- **Improvements with this approach**
 - **JeOS: drill down entire OS install**
 - **means more streamlined installation**
 - **Defined as a Open Virtualization Format (OVF) VM deployment file**
 - **Can be used under Xen or Virtual box**

ESG Data Node Stress Test

- Our ESG data node now publishes to the PCMDI gateway
- CSIRO's CMIP3 data products being re-published on the ESG node at the NF
- Acts as a stress test to identify
 - JVM overheads
 - Impact on the I/O subsystems
 - Running OProfile over the entire VM
- Can compare results to other international fast data transfers



National Computational Infrastructure

NCI National Facility

Supported by:

➤ Australian National University

➤ CSIRO, inc CAWCR

➤ Research Intensive Universities (Go8)

➤ Aust. Bureau of Met
➤ Geoscience Australia



An Australian Government Initiative

**National Collaborative Research
Infrastructure Strategy**

Contacts:

Ben.Evans@nf.nci.org.au

Joseph.Antony@nf.nci.org.au

<http://nf.nci.org.au>